Buyer to Seller Recommendation under Constraints

Cheng Chen
University of Victoria
Victoria, Canada
cchenv@uvic.ca

Lan Zheng
University of Victoria
Victoria, Canada
lanzheng@uvic.ca

Venkatesh Srinivasan
University of Victoria
Victoria, Canada
venkat@cs.uvic.ca

Alex Thomo
University of Victoria
Victoria, Canada
thomo@cs.uvic.ca

Kui Wu
University of Victoria
Victoria, Canada
wkui@uvic.ca

Anthony Sukow
Terapeak
Victoria, Canada
anthony@terapeak.com

arXiv:1406.0455v3 [cs.SI] 13 Jun 2014

ABSTRACT

The majority of recommender systems are designed to recommend items (such as movies and products) to users. We focus on the problem of recommending buyers to sellers which comes with new challenges: (1) constraints on the number of recommendations buyers are part of before they become overwhelmed, (2) constraints on the number of recommendations sellers receive within their budget, and (3) constraints on the set of buyers that sellers want to receive (e.g., no more than two people from the same household). We propose the following critical problems of recommending buyers to sellers: Constrained Recommendation (C-REC) capturing the first two challenges, and Conflict-Aware Constrained Recommendation (CAC-REC) capturing all three challenges at the same time. We show that C-REC can be modeled using linear programming and can be efficiently solved using modern solvers. On the other hand, we show that CAC-REC is NP-hard. We propose two approximate algorithms to solve CAC-REC and show that they achieve close to optimal solutions via comprehensive experiments using real-world datasets.

Categories and Subject Descriptors

H.3 [**Information Storage and Retrieval**]: Information Search and Retrieval—*Information filtering*; F.2.1 [**Analysis of Algorithms and Problem Complexity**]: Numerical Algorithms and Problems—*computations on matrices*

General Terms

Algorithms

Keywords

Buyer recommendation, optimization, approximation

1. INTRODUCTION

We study the problem of recommending buyers to sellers (RBS) in an online e-commerce site, such as eBay. This problem comes with its own set of challenges, which are quite different from those of recommending items to users. The first challenge is that we do not want to recommend the top buyers (in terms of their buying history) to all the sellers because those top buyers will be inundated by thousands of advertisement messages from eager sellers trying to reach out to them. Therefore we need constraints in place that limit the number of sellers each buyer can be recommended to. The second challenge is that sellers are normally under budget limitations, thus do not want more buyers recommended to them than what they can pay for. The third challenge is that we need to recommend top buyers to top sellers, middle-tier buyers to middle-tier sellers, and so on, as this maximizes the chance of making both the buyers and sellers happy; typically, top buyers buy a lot of merchandise, and it is the top sellers that have the best chance of satisfying their appetite. In general we would like to avoid top buyers being recommended to less than stellar sellers unless the latter have already been saturated with recommended buyers.

We do not normally have these challenges when recommending items (e.g., movies) to users (RIU). For example there is absolutely no problem recommending the best movies to all the users of an online site; on the contrary, it is widely applied. Also, there is no problem of recommending too many movies to one person as long as they are sorted by the rating prediction. And, finally, movie recommendation is a very "democratic" process; top movies are recommended all the time to less than stellar users, i.e., those that do not watch and rate many movies (cold-start users).

From all the above, it is clear that RBS is very different from RIU, and therefore, we need a new formalization and new algorithms to address RBS. To the best of our knowledge, we are the first to directly approach this important problem. This is because classical advertising is done in "broadcast" mode via TV commercials and sponsored links in search results, which do not have the aforementioned challenges. Another way of advertising is via targeted campaigns. This is typically done (or initiated) by a single merchant company and involves identifying the most promising customers to target for its own products. Again, such a process is different from RBS, where there are many sellers competing for the buyers.

In this paper, we introduce two central problems, Constrained Recommendation (C-REC) and Conflict-Aware Constrained Recommendation (CAC-REC), to address different RBS scenarios. We view RBS from the point of view of an online company, like eBay. Such a company typically associates a profit weight for each recommendation of a buyer to a seller. For a given seller, the more highly ranked the buyer, the more profitable the recommendation.

In C-REC, we assume that each buyer and seller has a specific "capacity." That is, a buyer cannot be recommended to more sellers than his capacity, and a seller cannot be

recommended more buyers than he has budget (capacity) for. Then the goal is to generate recommendations maximizing the total profitability (sum of recommendation weights), while respecting the capacity constraints of the buyers and sellers.

In CAC-REC, we extend C-REC in a natural way. Namely, we assume the buyers could be "in confict" with other buyers. For example, a buyer might use the same IP address or browser cookie as another buyer, which would indicate that these two buyers share the same computer and probably live together in the same household. A seller might not want to be recommended buyers that are in conflict with each other, as for example in the case when the seller sells household items such as TVs, washers, etc.

We show that C-REC can be modeled using linear programming and can be efficiently solved using modern solvers. On the other hand, we show that CAC-REC is NP-hard. We propose two approximate algorithms to solve CAC-REC and show that they achieve close to optimal solutions via comprehensive experiments using real-world datasets.

More specifically, our contributions are as follows:

1. We initiate the study of natural problems arising in RBS systems: How do we recommend high-value buyers to sellers under various constraints?

2. In the presence of capacity constraints (C-REC), we model the problem as integer linear programming and show that the problem can be optimally solved using LP solvers (i.e., the LP solutions are integral).

3. For the case of capacity and conflict constraints (CAC-REC), we prove that the problem is NP-hard and model it as semidefinite programming and integer linear programming. We present a greedy algorithm that is scalable and close to optimal.

4. We provide an extensive experimental evaluation on real-world datasets validating the claims of scalability and optimality made above.

2. RELATED WORK

RBS is related to the task of constraint-based recommendation (cf. [2, 16, 3, 4, 12, 15, 8]).

The first two works consider constraints on item features (attributes), thus they do not study the same type of constraints as we do. The rest of works are more closely related to our first, C-REC, problem.

Karimzadehgan et. al in [5, 3, 4] study the problem of optimizing the review assignments of scientific papers. Similarly to C-REC, they employ constraints on the quota of papers each reviewer is assigned. However, differently from C-REC, in their optimization setup, matching of reviewers with a paper is done based on matching of multiple aspects of expertise. Taylor in [12] also considers the paper review assignment problem. The difference from C-REC is that it does not consider an ordering on the reviewers and papers.

Xie, Lakshmanan, and Wood in [15] study the problem of composite recommendations, where each recommendation comprises a set of items. They also consider constraints including the number of items that can be recommended to a user. Their objective, however, is to minimize the cost of a recommended set of items when each item has a price to be paid.

Parameswaran, Venetis, and Garcia-Molina in [8] study the problem of course recommendations with course requirement constraints. Similarly as [15], the goal of [8] is to come up with set recommendations. However, the challenge they address is the modeling of complex academic requirements (e.g., take 2 out of a set of 5 math courses to meet the degree requirement). Such constraints are different from those that we consider.

To the best of our knowledge, the second problem we consider, *CAC-REC*, is completely new.

3. CONSTRAINED RECOMMENDATION (C-REC)

We first study the recommendation problem without considering conflict between buyers. To describe the problem formally, we model the buyer-seller network as a undirected, bipartite graph (Figure 1) $G = \langle (B, S), E, W \rangle$, where $B = \{b_1, b_2, \ldots, b_m\}$ denotes the list of buyers, $S = \{s_1, s_2, \ldots, s_n\}$ the list of sellers, $E \subseteq B \times S$ the set of edges, and $W : E \to \mathbb{R}^+$ the weights of the edges.

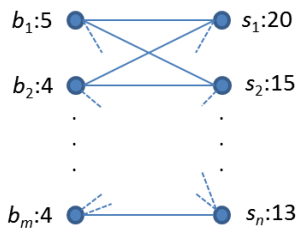

Figure 1: C-REC. The numbers are degree constraints.

For convenience, we slightly abuse the notation and use an mn-dimensional vector to denote $E = [e_{ij}]$, with $e_{ij} = 1$ indicating that there is an edge between buyer i and seller j and $e_{ij} = 0$ otherwise. Similarly, we use a vector to denote $W = [w_{ij}]$. If $e_{ij} = 0$, then $w_{ij} = 0$. When the subscripts are hard to read, as in the next section, we also use $W(i, j)$ to denote w_{ij}.

The weight of the edge between a buyer $i \in B$ and a seller $j \in S$, w_{ij}, reflects the profitability value if the buyer is recommended to the seller. In practice, we may pre-process the bipartite graph based on various business models. For instance, we may order the buyers or sellers based on money spent and earned, recency of buys and sells, etc. In addition, we may constrain the edges from a buyer to a ranked range of sellers, so that the buyer is not recommended to sellers well outside of her "tier".

It is often desirable that only a certain number of buyers is recommended to a seller, since a seller may be overwhelmed otherwise or because a seller needs to pay for the recommendation. Similarly, it is not reasonable to recommend a buyer to a large number of sellers, as the buyer might be annoyed later by many unsolicited requests. To avoid the problem, a good recommendation system should allow us to constrain the number of recommendations associated with individual buyers and sellers. In other words, we need to put degree constraints $D : B \cup S \to \mathbb{N}$ in the bipartite graph. We represent D with a $(m + n)$-dimensional column vector $D = [D(i)]^T$.

Denote $X = [x_{ij}]^T$ as the mn-dimensional column vector of 0-1 variables, with $x_{ij} = 1$ indicating buyer i is recommended to seller j and $x_{ij} = 0$ otherwise. The constrained recommendation (C-REC) problem is to *find the set of recommendations such that the total profit of the recommendation is maximized under the degree constraints*, i.e.,

$$\begin{aligned} \max_{X} \quad & WX \\ s.t. \quad & \mathbb{A}X(i) \leq D(i), \forall i, 1 \leq i \leq m+n \\ & x_{ij} \in \{0,1\}, \forall i,j, 1 \leq i \leq m, 1 \leq j \leq n, \end{aligned} \tag{1}$$

where matrix $\mathbb{A}$ is an $(m+n) \times mn$ matrix defined by

$$\underbrace{\begin{bmatrix} [e_{11},\ldots,e_{1n}] & & & \\ & [e_{21},\ldots,e_{2n}] & & \\ & & \ddots & [e_{m1},\ldots,e_{mn}] \\ [e_{11},0,\ldots,0] & & \ldots & [e_{m1},0,\ldots,0] \\ \ddots & & \ddots & \ddots \\ [0,\ldots,0,e_{1n}] & & \ldots & [0,\ldots,0,e_{mn}] \end{bmatrix}}_{(m+n)\times mn}. \tag{2}$$

The degree constraints are given by $\mathbb{A}X(i) \leq D(i)$, where $\mathbb{A}X(i)$ denotes the i-th element in (vector) $\mathbb{A}X$ and $D(i)$ the i-th element in D.

From the above formulation, it is easy to see that the C-REC problem could be reduced to the transportation problem in operations research [1, 10, 11], and as such we obtain an LP problem that can be solved efficiently by the modern solvers. Furthermore, it can be shown the LP solution is also integral.

4. CONFLICT-AWARE CONSTRAINED RECOMMENDATION (CAC-REC)

We now consider a natural generalization of problem C-REC. In some scenarios, sellers will prefer a *diverse* list of recommended buyers that avoids certain redundancies. For example, a seller might prefer that their list does not include more than one recommended buyer from each household. Advertising to more than one potential buyer in a household for a given merchandise, in most cases, is unnecessary. We will represent presence of such dependencies or conflicts between two buyers using (unweighted) conflict edges (Figure 2). We call this recommendation problem conflict-aware constrained recommendation (CAC-REC). The goal is to *compute a maximum weight subgraph satisfying the degree constraints as in C-REC with the additional requirement that the number of conflict edges within a list of buyers recommended to any particular seller is smaller than a threshold.*

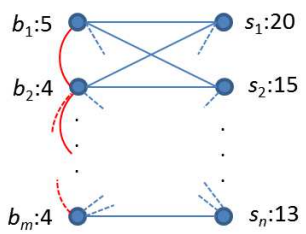

Figure 2: CAC-REC. The numbers are degree constraints and the red edges represent conflicts.

To describe CAC-REC more precisely, the input to CAC-REC consists of the following information:

1. An undirected, weighted graph $G = \langle (B,S), E \cup C, W \rangle$ with $E \subseteq B \times S$, $C \subseteq B \times B$ and weights $W : E \to \mathbb{R}^+$;
2. Degree constraints $D : B \cup S \to \mathbb{N}$;
3. A conflict threshold t.

The goal in CAC-REC is to compute a maximum weight subgraph G' of G, $G' = ((B,S), E' \cup C', W)$, that satisfies the following two constraints:

1. For any i in $B \cup S$, $d_{G'}(i) \leq D(i)$.
2. For any k in S, $|\{(i,j)|(i,k),(j,k) \in E', (i,j) \in C'\}| \leq t$.

Here, $d_{G'}(i)$ denoted the degree of vertex i in the subgraph G'.

We will next show that considering the conflict between buyers increases the complexity of the recommendation problem significantly.

4.1 NP-Hardness Result for CAC-REC

In this section, we provide strong evidence that CAC-REC is highly unlikely to have an efficient (i.e, polynomial time) algorithm by showing that it is NP-hard.

THEOREM 1. CAC-REC *is NP-hard.*

PROOF. We give a polynomial-time reduction from the NP-hard problem REVENUE MAXIMIZATION IN INTERVAL SCHEDULING.

REVENUE MAXIMIZATION IN INTERVAL SCHEDULING (RMIS)

Instance: A set $M = \{m_1, m_2, \ldots, m_t\}$ of t machines and a set $J = \{j_1, j_2, \ldots, j_n\}$ of n jobs. For each job j in J, we are given three parameters: (1) $S(j)$, the set of machines on which j can be processed (2) $R(j,-)$, the set of possible revenues obtained when job j is processed on different machines (3) $I(j)$, the time interval during which job j must be processed.

Goal: Find a feasible schedule of a subset of jobs on the machines that maximizes the total revenue of the jobs scheduled.

We now describe a reduction from RMIS to CAC-REC. Given an instance I of the revenue maximization problem, construct a graph $G(I)$, which is an instance of CAC-REC. Define $G(I) = \langle (J,M), E \cup C, W \rangle$ with $E \subseteq J \times M$, $C \subseteq J \times J$ and weights $W : E \to \mathbb{R}^+$ as follows.

- $E = \{(j_k, m_l) | m_l \in S(j_k)\}$.
- $C = \{(j_k, j_l) | I(j_k) \cap I(j_l) \neq \emptyset\}$.
- $W(j_k, m_l) = R(j_k, m_l)$.
- $D(j_k) = 1$ for all k and $D(m_l) = n$ for all l.
- $t = 0$.

We now explain the reduction above. There is an edge between job j_k and machine m_l if m_l belongs to $S(j_k)$, the set of machines in which job j_k can be processed. There is a conflict edge between job j_k and job j_l if their time interval for processing overlap. The weights on an edge (j_k, m_l) represent the revenue obtained if job j_k is processed on machine m_l. Since each job can be assigned to at most one machine, their degree constraints are set to 1. There is no constraint on the number of jobs assigned to any machine and hence their degree constraint is set to n. Finally, there must be no conflict between two jobs assigned to same machine. Therefore, t is set to 0.

It can be easily seen that an optimal solution for $G(I)$, an instance of CAC-REC yields an optimal solution for I, an instance of RMIS. In other words, a maximum weight subgraph of $G(I)$ satisfying the degree constraints and conflict constraints as described above exactly corresponds to a revenue maximizing schedule in I. Furthermore, we observe that the reduction above is a polynomial-time reduction. Therefore we conclude that CAC-REC is NP-hard. □

4.2 A SDP Algorithm for CAC-REC

In the previous section, we showed that CAC-REC is NP-hard. In this section and next, we will design two efficient algorithms for CAC-REC that provide high-quality solutions that are close to optimal.

Our first algorithm for CAC-REC is based on a semidefinite programming based approach. To understand the motivation for this approach, recall that we described a LP formulation for C-REC in Section 3. Using the terminology from Section 3 and 4.1, the conflict constraint can be described as follows:

$$\sum_{(j_k,j_l)\in C} x_{ki}x_{li} \leq t \;\forall\; i \in S \tag{3}$$

That is, the conflict constraint is quadratic. We use a single t for illustration purpose. In practice, different sellers can have different values of t. We will now show how to formulate CAC-REC as a semidefinite program. Define a $mn \times mn$ symmetric matrix $\mathbb{Y} = XX^T$ where X is as in Section 3. The CAC-REC problem can be described as

$$\begin{aligned} \max \quad & Trace(\mathbb{W}\mathbb{Y}) \\ s.t. \quad & Trace(\mathbb{D}_i^b\mathbb{Y}) \leq D(i), \forall\; i \in B \\ & Trace(\mathbb{D}_i^s\mathbb{Y}) \leq D(i), \forall\; i \in S \\ & Trace(\mathbb{C}_i\mathbb{Y}) \leq t, \forall\; i \in S \\ & \mathbb{Y} = XX^T \succeq 0 \end{aligned} \tag{4}$$

where $\mathbb{W}, \mathbb{D}_i^b, \mathbb{D}_i^s$ and $\mathbb{C}_i$ are suitably defined $mn \times mn$ symmetric matrices described below.

1. $\mathbb{W}$ is a diagonal matrix with diagonal weights w_{ij}.
2. $\mathbb{D}_i^b$ is diagonal matrix with a 1 for row indexed by (i,j) if $(i,j) \in E$ and 0 otherwise.
3. $\mathbb{D}_i^s$ is diagonal matrix with a 1 for row indexed by (k,i) if $(k,i) \in E$ and 0 otherwise.
4. Finally, $\mathbb{C}_i$ is a matrix with entries 1/2 and 0. An entry indexed by row (j,i) and column (k,i) is equal to 1/2 if $(j,i),(k,i) \in E$ and $(j,k) \in C$. It is 0 otherwise.

Our SDP based algorithm for CAC-REC is as follows:

1. Solve the semidefinite program relaxation to obtain optimal solution $\mathbb{Y}$.
2. Using the Cholesky decomposition [9] of $\mathbb{Y}$, obtain the vectors x_{ij} corresponding to $\mathbb{Y}$.
3. Use a two-step rounding procedure, random projection followed by threshold rounding, to obtain $0, 1$ values for x_{ij}.

In step (1) of our SDP algorithm, the SDP described above is solved using a generic SDP solver. The output of step (1) is a semidefinite matrix $\mathbb{Y}$. In step (2) of our algorithm, we use a well-known fact that any semidefinite matrix $\mathbb{Y}$ can be written as $\mathbb{Y} = \mathbb{V}\mathbb{V}^T$ where $\mathbb{V}$ is a $mn \times mn$ lower triangular matrix. This decompostion is known as Cholesky decomposition of $\mathbb{Y}$ [9]. The columns of $\mathbb{V}$ give us a vector solution for the variables x_{ij}. Thus, the output of step (2) of our algorithm are mn vectors, one for each x_{ij}. These vectors correspond to the optimal solution of the SDP. In the last step of our algorithm, we convert the vectors x_{ij} to integral $\{0,1\}$ values using a two-step rounding procedure. In the first step, we convert x_{ij}'s to fractional values by a random projection. That is, we pick a random vector x of dimension mn by picking each of its coordinates from the normal distribution $N(0,1)$ and define each $x_{i,j}$ as the length of its projection on to x. Finally, we sort x_{ij} and round each non-zero value to 1 provided doing so does not violate the degree constraints or the conflict constraints. Otherwise, we set it to 0.

4.3 ILP Formulation of CAC-REC

Solving the SDP formulation requires large enough physical memory to store the $mn \times mn$ matrix $\mathbb{Y}$. In practice (e.g., the eBay purchase graph of a certain category), large values for m (the number of buyers) and n (the number of sellers) inevitably restrict the applicability of the SDP approach. This limitation, however, can be alleviated if we could model CAC-REC as an integer linear programming (ILP) problem.

In order to achieve this goal, we introduce a new 0-1 variable $z_{i,(j,k)}$ to formulate Inequality 3 as a linear constraint. For each seller i, $z_{i,(k,l)}$ equals 1 if and only if there is a conflict edge between two buyers k and l, and both edge e_{ki} and e_{li} are recommended in the graph. Using the terminology from Section 3 and 4.1, this constraint can be described as follows:

$$(1 - x_{ki}) + (1 - x_{li}) + z_{i,(k,l)} \geq 1 \;\forall\; i \in S, \quad \forall k,l \in B \tag{5}$$

$$x_{ki} + x_{li} - 2z_{i,(k,l)} \geq 0 \;\forall\; i \in S, \quad \forall k,l \in B \tag{6}$$

$$\sum_{(k,l)\in C_i} z_{i,(k,l)} \leq t \;\forall\; i \in S \tag{7}$$

In Constraint (6), C_i is defined as follows: $C_i = \{(k,l) \in C | (k,i) \in E \wedge (l,i) \in E\}$.

The linear conflict constraints can be easily incorporated into C-REC (Problem 1). Let c denote the total number of conflict constraints with respect to all sellers in the graph. Then we can obtain a linear programming formulation, where $X = [x_{ij}]^T$ is a $(mn + c)$-dimensional column vector of 0-1 variables. In addition, matrix $\mathbb{A}$ and vectors W and D can be changed accordingly.

By eliminated the need to store large dimensional matrices, now we can use a ILP solver to tackle CAC-REC problems of larger sizes. As opposed to C-REC, obtaining an integer solution in CAC-REC is NP-hard. In order to further improve efficiency, we use a rounding procedure after solving the linear program relaxation. Our LP based algorithm for CAC-REC is as follows:

1. Solve the linear program relaxation to obtain optimal solution X.
2. Sort the first mn elements of X from largest to smallest. We round each non-zero value to 1 provided doing so does not violate the degree constraints or the conflict constraints. Otherwise, we set it to 0.

4.4 A Greedy Algorithm for CAC-REC

In this section, we describe and study the performance of a simple greedy algorithm for this problem. This algorithm has the advantage that it is highly scalable and provides good quality solutions in practice.

The greedy algorithm, denoted as GREEDY, for CAC-REC is as follows:

1. Sort all the edges in E by weight from largest to smallest.
2. To construct the maximum weight subgraph G', consider every edge in the sorted list. Add this edge to G' if doing so does not violate any degree constraint or conflict constraint.
3. Continue until we reach the end of the sorted list.

We will now prove a theoretical guarantee on the performance of GREEDY.

THEOREM 2. *Let $d = \max_{v \in B} |\{(v, v')|(v, v') \in C\}|$. Algorithm GREEDY is a $(2 + d)$- approximation algorithm.*

PROOF. We use the concept of a *k-extendible system* to provide performance guarantees of a greedy algorithm. Mestre introduced the notion of a k-extendible system in his study of the performance of the greedy technique as an approximation algorithm [7].

DEFINITION 1 (k-EXTENDIBLE SYSTEM [7]). *Let U be a finite set and $\mathcal{F}$, $\mathcal{F} \subseteq 2^U$, be a collection of subsets of U. Set system $(U, \mathcal{F})$ is called a* k-extendible system *if it satisfies the following properties:*

1. Downward-closure*: If $A \subseteq B$ and $B \in \mathcal{F}$, then $A \in \mathcal{F}$.*
2. Exchange*: Let $A, B \in \mathcal{F}$ with $A \subseteq B$, and let $x \in U - B$ be such that $A \cup \{x\} \in \mathcal{F}$. Then there exists $Y \subseteq B - A$, $|Y| \leq k$, such that $(B - Y) \cup \{x\} \in \mathcal{F}$. In other words, let us start with any choice of two sets A and B such that B is an extension of A. Suppose that there is an element x such that the set A with x added to it also belongs to $\mathcal{F}$. Then we will be able to find a subset Y inside B of size at most k such that if we remove the elements of Y from B and add the element x to the resulting set, it will also belong to the collection $\mathcal{F}$.*

Informally, Mestre showed that if the set of all feasible solutions forms a k-extendible system, algorithm GREEDY gives a k-approximation algorithm. That is, on any instance, the solution output by GREEDY differs from the optimal solution by a multiplicative factor of at most k. We now state his result more formally.

THEOREM 3 (MESTRE [7]). *Let $(U, \mathcal{F})$ be a k-extendible system for some k. Let $W : U \to \mathbb{R}^+$ be a positive weight function on U. Then, the greedy algorithm gives a k-approximation algorithm for the optimization problem that asks to determine $\max_{F \in \mathcal{F}} W(F)$ where $W(F) = \sum_{s \in F} W(s)$ for any $F \in \mathcal{F}$.*

To apply this result to our problem, we will check that the set of all feasible solutions to CAC-REC forms a $(2+d)$-extendible system. For the CAC-REC problem, $U = E$ and $\mathcal{F}$ be the set of all subgraphs of G satisfying the degree and conflict constraints. Then it is easy to see that $(U, \mathcal{F})$ is downward closed. That is, removing an edge from a feasible solution H will always result in be a feasible solution as this will not cause any violation of constraints.

For the exchange property, consider the case when a new edge $e = (u, v)$, $u \in B$ and $v \in S$, is added to a feasible solution H. We make two observations: (1) Adding e could result in violation of degree constraint at u and v. However, this can be rectified by removing two other edges, one incident on u and other incident on v; (2) Adding e could result in a violation of the conflict constraint at v. Rectifying this could require removing at most d edges where $d = \max_{v \in B} |\{(v, v')|(v, v') \in C\}|$. Therefore, we obtain a $(2 + d)$-extendible system. □

We remark that this analysis is worst-case. In practice, GREEDY shows far superior performance, as demonstrated in our later test with real-world dataset.

5. EXPERIMENTAL EVALUATION

In order to illustrate the proposed algorithms' optimality and scalability, we conducted comprehensive experiments on real-world datasets for *C-REC* and *CAC-REC*. The datasets (provided by Terapeak[1]) include three-month transaction records across all categories of eBay Canada in 2013. We note that even though we used data from Terapeak, similar datasets can be easily collected via the eBay API. We use the sales data of a specific category (Cell phones and Accessories) in our study, which contains 18,742 buyers and 1,884 sellers after filtering out buyers who only purchased from a single seller and sellers who sold items to no more than 5 distinct buyers. Then we sort buyers, as well as sellers, by the total monetary purchases and total profits, respectively. The weight on each edge comes from the real-world sales data and will be explained in more detail in the following experiments.

All experiments were run on a 64-bit Ubuntu 12.04 desktop of 3.40GHz * 8 Intel Core i7 CPU and 3.8 GB memory.

5.1 C-REC

We use a fast linear programming (LP) solver, Gurobi[2] for Matlab, to solve *C-REC* at different scales. The purpose is to observe the scalability of the LP approach.

[1]Terapeak is an E-commerce company, helping sellers on eBay or Amazon to measure and boost their sales performance. http://www.terapeak.ca/

[2]http://www.gurobi.com/

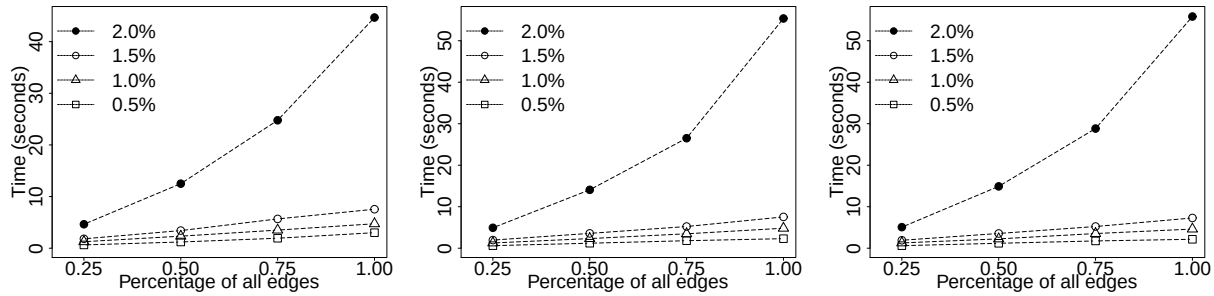

(a) Degree Constraint Ratio: 10% (b) Degree Constraint Ratio: 20% (c) Degree Constraint Ratio: 30%

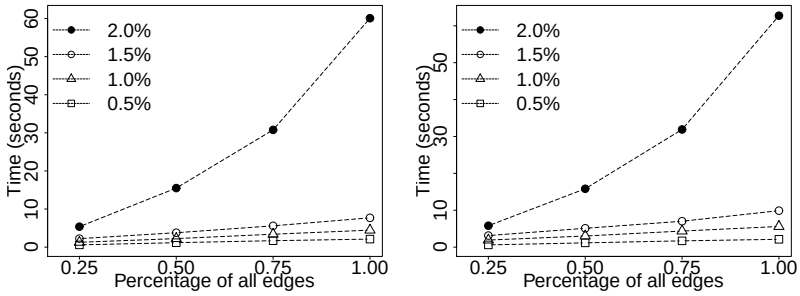

(d) Degree Constraint Ratio: 40% (e) Degree Constraint Ratio: 50%

Figure 3: C-REC running time of different degree constraint ratios. Similar trends are seen for running time in all scenarios, i.e., C-REC with a higher density requires a longer time to be solved and the running time increases when the size of the graph is scaled up.

Adding edges between sorted buyers and sellers, we create four bipartite graphs of different density settings, roughly at 0.5%, 1.0%, 1.5% and 2.0%. Edges are generated in a manner that each seller has the same number of connected buyers. For example, if the density is 0.5%, the first seller (top ranked) is connected with the first 90 buyers (from 1 to 90), and the second seller is connected with buyers from 11 to 100, and so on. We also extract three subsets of different sizes for each bipartite graph, i.e., using 25%, 50% and 75% of the total number of edges (the number of buyer and seller nodes decreases accordingly). The weight of an edge is the sum of the buyer's total monetary purchases (on all sellers within the category) and the seller's total profits (from all buyers within the category). For each buyer (seller) in the bipartite graph, the degree constraint ratio is the proportion of the maximum number of recommended sellers (buyers) among all candidates. We choose constraint ratios from the set $\{0.1, 0.2, 0.3, 0.4, 0.5\}$ to illustrate their impact on the running time. Every node (buyer or seller) in the experiment has the same degree constraint ratio.

Figure 3 shows the running time of different experimental settings. The four curves in each subfigure correspond to results of bipartite graphs with different densities, i.e., 0.5%, 1.0%, 1.5%, and 2.0%, respectively. The final result is computed as the average of five runs. We make the following observations:

1. The curve of a higher density is always positioned above the one of a lower density, indicating that it takes a longer time to solve a *C-REC* of a higher density. Specifically, the 2.0% density *C-REC* requires much more time compared to others of smaller density.
2. The running time increases when we scale up the size of the bipartite graph. For example, the rising curve of the 2.0% density becomes steeper as we add more edges. We observed a similar trend in other density curves, but they are not as obvious as the 2.0% one.
3. The running time slightly changes when the degree constraint ratio becomes larger. For example, the running times of the 2.0% curve of all edges for each degree constraint ratio (subfigures (a), (b), (c), (d), and (e)) are 44.67, 55.36, 55.83, 60.08 and 62.74 seconds, respectively.

From Figure 3, we observe that no matter how we change the size of the bipartite graph and degree constraints, it only takes about one minute to solve the largest instance of *C-REC*. The results suggest that the LP approach for *C-REC* is highly scalable.

5.2 CAC-REC

Now we introduce conflict constraints between pairs of buyers.

5.2.1 SDP Formulation of CAC-REC

As discussed in Section 4, the SDP approach of *CAC-REC* can be approximately solved by a SDP solver plus a rounding procedure or by a greedy algorithm. We use SDPT3 [13, 14] as the SDP solver and we write a Python script to implement the GREEDY algorithm. We compare the solution of each method, optimal SDP (solution without rounding), rounded SDP, and GREEDY.

The applicability of the SDP approach for *CAC-REC* severely suffers from the limitation of physical memory, i.e., the need

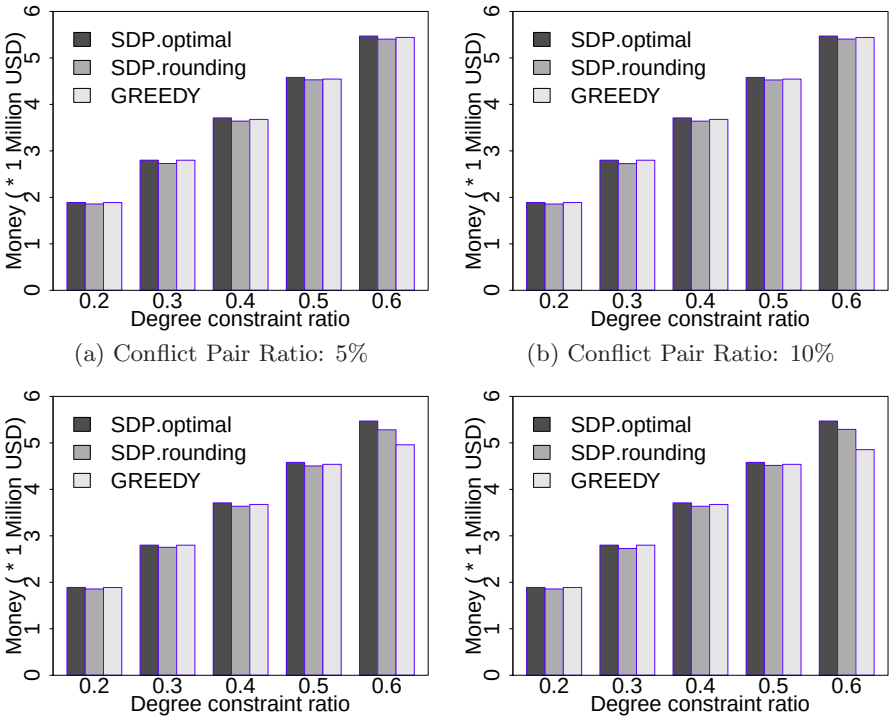

(a) Conflict Pair Ratio: 5%

(b) Conflict Pair Ratio: 10%

(c) Conflict Pair Ratio: 15%

(d) Conflict Pair Ratio: 20%

Figure 4: CAC-REC money solution of different conflict pair ratios. When degree constraint ratio and conflict pair ratio are both low, GREEDY shows close to optimal solutions while SDP with rounding is weaker.

to store a large-dimensional matrix [6]. One way to alleviate the problem is to factorize the bipartite graph into several independent subgraphs, which do not share common buyers or sellers. Then we can run the SDP solver on each subgraph individually and merge all results. This can be achieved considering the way we create the edges between buyers and sellers, i.e., each seller is connected to buyers in a sorted subsequence. In the following, we show results obtained on a small subgraph to avoid repeated comparison.

We create a subgraph of small size consisting of top 5 sellers and top 26 buyers (ranked by monetary). Each seller is connected with 10 buyers so that each buyer can be assigned to multiple sellers. We use two types of edge weight, money and node rank. Money weight is computed in the same way as in the experiment of *C-REC*, i.e., the sum of the buyer's total monetary purchases (on all sellers within the category) and the seller's total profits (from all buyers within the category). Of course, any monotone weight function can be used. Rank weight equals the multiplication of a big constant (the total number of buyer and seller nodes in the entire graph, $20,626$ in our case) and the reciprocal of the sum of the buyer's rank and the seller's rank.

We create different number of conflict buyer pairs by randomly sampling the number of buyer pairs by the following percentages of total possible buyer pairs, $\{5\%, 10\%, 15\%, 20\%\}$. Similar to *C-REC*, we also set different degree constraint ratios for each node, $\{20\%, 30\%, 40\%, 50\%, 60\%\}$. In addition, we use a constant ("1") for each seller's conflict constraint, which means each seller can at most accept one conflict buyer pair. We use this constant to amplify the impact of conflict constraint on the small-scale subgraph. For SDP with rounding, we run the randomized rounding procedure 20 times and take the best solution.

Figure 4 and Figure 5 depict the comparison of three approaches for the money weight and rank weight, respectively.

In Figure 4, we observe that regardless of conflict pair ratio, solutions of the three methods are very similar when degree constraint ratio is smaller than 60%, with SDP with rounding being slightly weaker than others. The reason is that when the degree constraint is tight, the conflict constraint of each seller is less likely to be activated, i.e., chances are rare for multiple conflict buyers to be recommended to a seller. The difference arises when the degree constraint ratio and the conflict pair ratio are both weak and the higher conflict pair ratio results in larger performance drop for both SDP with rounding and GREEDY (e.g., comparing the sets of bars of 60% degree constraint ratio in subfigures (c) and (d)).

Figure 5 shows a similar performance change trend for the three approaches. For example, the solution difference becomes larger when degree constraints are weaker and the performance of SDP with rounding and GREEDY gradually decreases as the conflict pair ratio increases. Meanwhile, we also observe that the performance of GREEDY is always superior to that of SDP with rounding.

Both SDP with rounding and GREEDY achieve close to optimal solutions. Specifically in the experiments, GREEDY exhibits a far superior performance compared to the theoretical analysis. In addition, the most important advantage of GREEDY is that it scales very well even when applied to large-scale datasets. We show its scalability in Figure 6. We use the same 2.0% density graphs of different size as in the

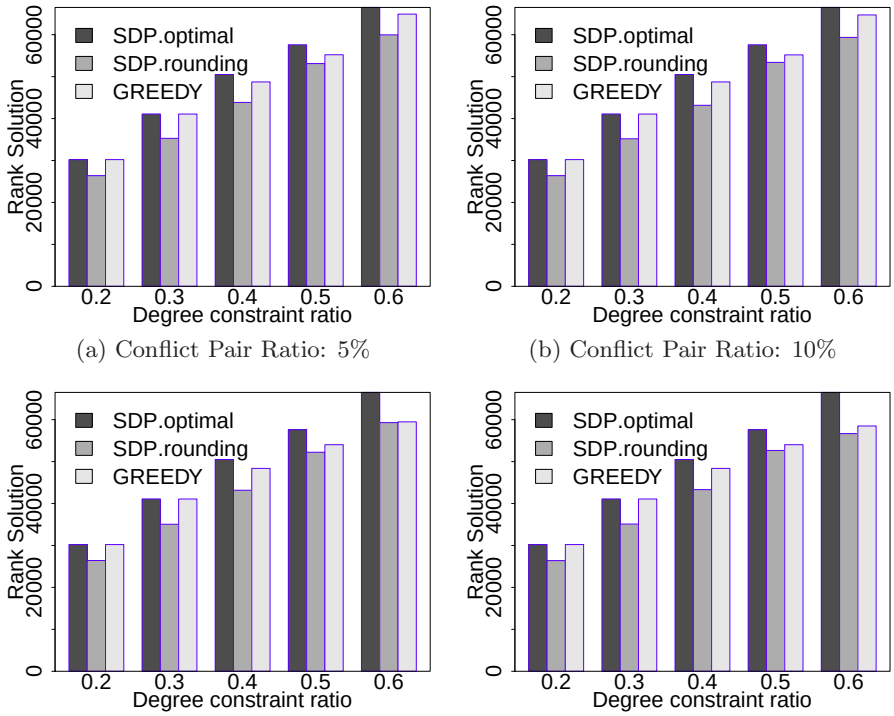

(a) Conflict Pair Ratio: 5%

(b) Conflict Pair Ratio: 10%

(c) Conflict Pair Ratio: 15%

(d) Conflict Pair Ratio: 20%

Figure 5: CAC-REC rank solution of different conflict pair ratios. GREEDY achieves close to optimal performance in all cases. The solutions obtained by SDP with rounding are slightly worse compared to GREEDY.

C-REC experiment. The settings of degree constraint ratio, conflict pair ratio are 60% and 10%, respectively. For each seller, the conflict threshold t (refer to Section 4) is set to be 50% of the total number of conflicting buyer pairs associated to the seller.

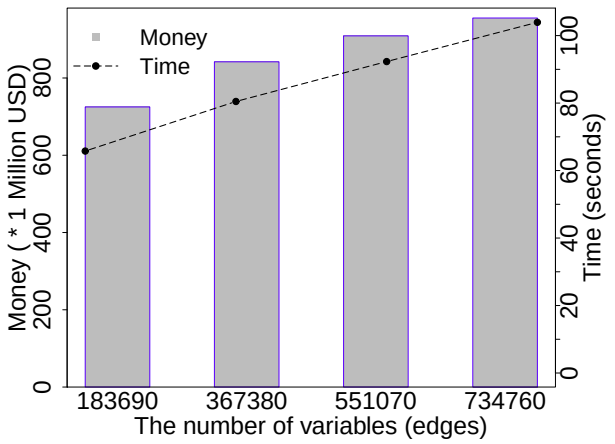

Figure 6: Greedy algorithm on large-scale datasets showing its scalability.

In Figure 6, the running time increases nearly linearly, and it only requires around 104 seconds to get a solution when the number of variables (edges) is considerably large, 734, 760!

5.2.2 *ILP Formulation of CAC-REC*

ILP formulation enables us to take full advantage of the LP solver (Gurobi) to solve CAC-REC problems with larger sizes. In this section, we perform larger-scale experiments to compare solutions of different methods, i.e., ILP, LP with rounding and GREEDY.

We create a 0.16%-density bipartite graph using the same method described in Section 5.1. The full graph consists of 18742 buyers, 1884 sellers and 56520 edges. We also extract three subsets of different sizes from the full graph, i.e., using 25%, 50% and 75% of the total number of edges (the number of buyer and seller nodes decreases accordingly). The degree constraint ratio, conflict pair ratio are 50% and 10%, respectively. The conflict threshold t (refer to Section 4) for each seller is set to be 50% of the total number of conflicting buyer pairs associated to the seller. Figure 7 shows the solution comparison of different methods on different datasets.

Figure 7 shows very promising results for LP with rounding and GREEDY algorithms on both money weights and rank weights; they are only slightly worse than the optimal integer solution obtained by ILP. Comparing to the SDP experiments, this experiment also verifies the effectiveness of both LP with rounding and GREEDY on large datasets because the graph we use in this section is considerably larger than the one used for SDP experiments. The density (0.16%) is also 10 times larger than the real-world graph (0.016%). Therefore, our LP formulation improves the scalability of solving larger-scale CAC-REC problems. In ad-

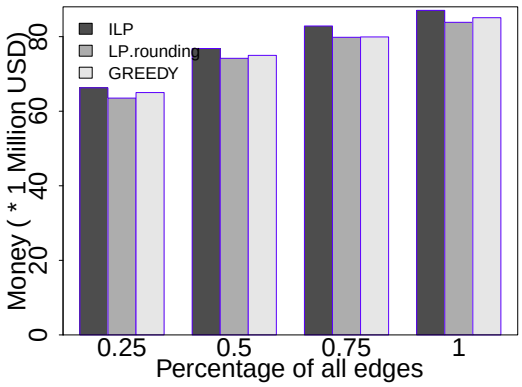

(a) Money Solution

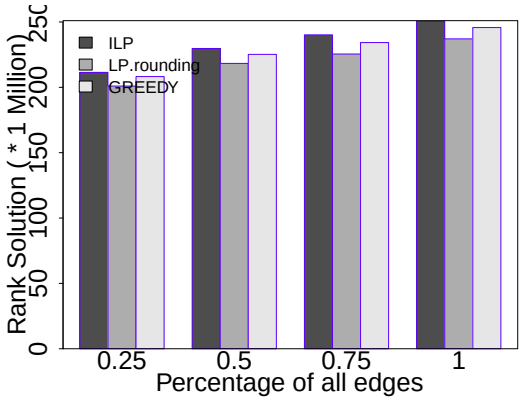

(b) Rank Solution

Figure 7: Linear Program experiments of CAC-REC on large-scale datasets.

dition, while ILP is NP-hard (running time may be long[3]), we can use the much faster GREEDY method or LP with rounding to improve the efficiency.

6. CONCLUSIONS

We introduced a novel recommendation problem that aims at recommending buyers to sellers (RBS) under capacity and conflict constraints. We provided formal definitions of two types of RBS, C-REC and CAC-REC, addressing different RBS scenarios. We showed that C-REC could be effectively solved using linear programming. By considering the conflict between buyers, however, the complexity of RBS increases significantly. We proved that CAC-REC is NP-hard. Then we proposed a SDP algorithm with a rounding procedure and a greedy algorithm to solve CAC-REC. Our results of extensive experiments using real-world datasets demonstrated that the proposed algorithms can achieve close to optimal solutions. Finally, we showed that the greedy algorithm is highly scalable.

Acknowledgments

We would like to thank Terapeak Inc. for providing the eBay data for this research.

[3]In our experiment in this section, the longest running time of ILP is approximately 10 minutes.

7. REFERENCES

[1] O. Amini, D. Peleg, S. Pérennes, I. Sau, and S. Saurabh. On the approximability of some degree-constrained subgraph problems. *Discrete Applied Mathematics*, 160(12):1661–1679, 2012.

[2] A. Felfernig and R. D. Burke. Constraint-based recommender systems: technologies and research issues. In *ICEC*, page 3, 2008.

[3] M. Karimzadehgan and C. Zhai. Constrained multi-aspect expertise matching for committee review assignment. In *CIKM*, pages 1697–1700, 2009.

[4] M. Karimzadehgan and C. Zhai. Integer linear programming for constrained multi-aspect committee review assignment. *Information processing & management*, 48(4):725–740, 2012.

[5] M. Karimzadehgan, C. Zhai, and G. Belford. Multi-aspect expertise matching for review assignment. In *CIKM*, pages 1113–1122. ACM, 2008.

[6] M. Kocvara and M. Stingl. On the solution of large-scale sdp problems by the modified barrier method using iterative solvers. *Math. Program.*, 109(2-3):413–444, 2007.

[7] J. Mestre. Greedy in approximation algorithms. In *ESA*, pages 528–539, 2006.

[8] A. G. Parameswaran, P. Venetis, and H. Garcia-Molina. Recommendation systems with complex constraints: A course recommendation perspective. *ACM Trans. Inf. Syst.*, 29(4):20, 2011.

[9] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical recipes in C: The art of scientific computing (second edition)*. Cambridge University Press, 1992.

[10] K. R. Rebman. Total unimodularity and the transportation problem: a generalization. *Linear Algebra and its Applications*, 8(1):11 – 24, 1974.

[11] A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1986.

[12] C. J. Taylor. On the optimal assignment of conference papers to reviewers. Technical Report MS-CIS-08-3, Computer and Information Science Department, University of Pennsylvania, 2008.

[13] K. C. Toh, M. Todd, and R. H. Tütüncü. Sdpt3 – a matlab software package for semidefinite programming. *Optimization Methods and Software*, 11:545–581, 1999.

[14] R. H. Tütüncü, K. C. Toh, and M. J. Todd. Solving semidefinite-quadratic-linear programs using sdpt3. *Math. Program.*, 95(2):189–217, 2003.

[15] M. Xie, L. V. S. Lakshmanan, and P. T. Wood. Breaking out of the box of recommendations: from items to packages. In *RecSys*, pages 151–158, 2010.

[16] M. Zanker and M. Jessenitschnig. Case-studies on exploiting explicit customer requirements in recommender systems. *User Model. User-Adapt. Interact.*, 19(1-2):133–166, 2009.